

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

**LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.**

www.theredhatsummit.com

Samba in Red Hat Enterprise Linux 6, clustering and other new Features

Simo Sorce
Principal Software Engineer
Red Hat, Inc.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Agenda

- New Features in Red Hat Enterprise Linux 6
 - Protocols, Interoperability, Management, ..
- Samba and Clustering
 - CTDB and how it works
- Samba Roadmaps
 - what is the community working on
 - Future developments



New features – Protocol support

- IPv6 support
 - Full support for single and dual stack setups. Ipv6-ification effort started in 2008 with samba 3.2
 - No additional configuration need, but IPv6 support relies on DNS (mDNS) for name resolution, no NETBIOS (this protocol is locked into IPv4).
- Support for the new SMB2 protocol
 - Currently (samba 3.5) considered experimental
 - Must be explicitly enabled (max protocol = smb2)
 - SMB2 is “better” but introduces new semantics



New Features - Interoperability

- Enhanced trusted domain support
 - Cross forest, transitive, one way
- Extended interoperability with Windows 7 and Windows 2008 R2
 - Supports all Windows security settings (sign/seal/etc)
 - Support for Windows 7 joined to a Samba DC
 - More information on:
<http://wiki.samba.org/index.php/Windows7>



New Features – Interoperability (cont.)

- Automatic machine password change as domain member
 - AD domain controllers may require domain members to change their password regularly
 - This requirement is addressed by a new option:
 - `machine password timeout = 604800` (default)
 - A value of 0 disables password changes
 - Running `winbindd` is required in order to use this feature
 - See “`wbinfo -c`” to force changes manually



New Features - Encryption

- Support for LDAP sign&seal policies
 - AD domains may require all LDAP traffic to be signed and/or sealed (integrity/privacy).
 - This requirement is addressed by a new option:
 - `client ldap sasl wrapping = plain` (default)
 - `client ldap sasl wrapping = sign`
 - `client ldap sasl wrapping = seal`
 - Sealing implies signing
- Encrypted SMB transport
 - NOTE: Windows clients do not support this extension
 - Requires posix extensions to be enabled



New Features - Management

- New, registry based, configuration mode
 - Samba can store the configuration in a registry file editable via network (`net conf` or `regedit32.exe`)
 - `HKLM\Software\Samba\smbconf`
 - Activated by `config backend = registry`
- Better support for native Windows management tools (`mmc`, `usrmgr`)
- New NetApi remote management api (like `netapi32.dll`)
- Graphical join utility tool (uses registry configuration)



Clustering: What kind of clustering ?

- Clients limitations determine what we can support
 - High availability, all nodes are active and provide the same data
 - Scalability for large number of clients, each client is redirected to a different node (DFS, IP failover)
 - Does not provide connection persistence on node failure
 - Clients cannot use more than one server at a time (no parallelization)
 - Clients must reconnect if a node fails, and re-establish state
- Goal: scale up clients with linear performance increase



Clustering: Samba + CTDB

- CTDB is a specialized distributed database used to clusterize Samba on top of a distributed file system like GFS2, GPFS, OCFS, ...
- By using CTDB Samba is able to properly share files from multiple nodes at the same time (active-active) maintaining full SMB semantics **without** degrading performances.
- CTDB also provides HA features such as node monitoring, node failover and IP takeover to simplify Samba cluster management, and can also be used to control other applications (like NFS sharing).

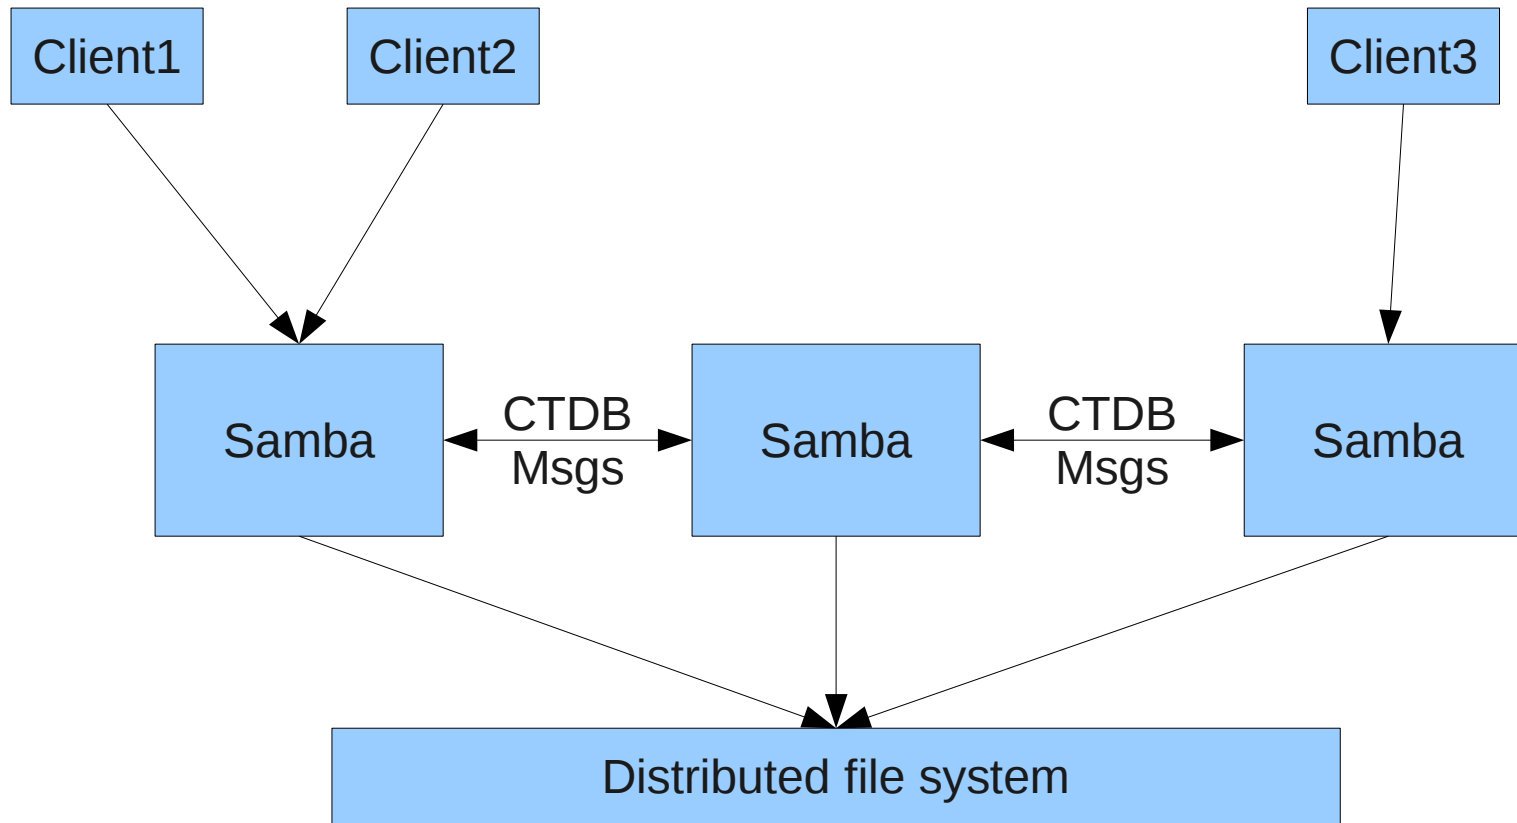


Why is the CTDB approach fundamental ?

- In order to fully support SMB semantics, all samba servers need to have very strict coordination.
- Classic sharing methods required to share samba's TDB databases over a distributed filesystem.
- Samba use locking very aggressively over TDB files
 - No distributed file system can cope with the amount of locking and data manipulation done on a traditional TDB database with adequate performances.
- CTDB allows to intelligently share data and messages between nodes only when it is really necessary.



Clustering



How does CTDB work ?

- CTDB uses *persistent* and *normal* databases.
- Data is always stored on a local, fast, file system.
 - Replication of *persistent* databases is performed by using a global transaction, and then distributing data to all the nodes when the transaction is complete.
 - Replication of *normal* databases is on-demand. Records are retrieved by a node only when needed, and cached locally.
- For *normal* databases every record belongs to a *data master* that has the authoritative copy of the record.



CTDB (cont.)

- The *data master* is dynamically determined based on usage in order to minimize inter-node communication
- When a node goes down a recovery procedure is performed by a designated *recovery master* that reclaims any orphaned record from the surviving nodes caches and becomes the new *data master* for those records.
- Normally the recovery procedure depends on holding a lock on a file on the underlying distributed file system. (lockless recovery is now available for testing)



CTDB and Samba configuration

- In a clustered environment it is very useful to be able to distribute configuration data automatically to every node.
- By using the registry configuration method, configuration can be immediately distributed to all nodes by CTDB as a persistent database
- Other databases containing users, id mappings, secrets are already shared through CTDB.



Samba Roadmaps

- Samba 3.5.x in RHEL 6
- New upstream release schedule changed to 9 months
- Next Samba release 3.6.0 expected by end of 2010
- SMB2 Unix Extensions to be designed together with Microsoft. Announced at SambaXP.
- Samba4 making great progress
 - Full AD replicating DC demonstrated at SambaXP
- Merge of current 3.x code and 4.x code expected next year leading to a 4.0.0 release (no firm schedule yet)



Q & A

Questions ?

Simo Sorce - ssorce@redhat.com

Samba - <http://www.samba.org>

CTDB - <http://ctdb.samba.org>

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



FOLLOW US ON TWITTER

www.twitter.com/redhatsummit

TWEET ABOUT IT

[#summitjbw](https://twitter.com/summitjbw)

READ THE BLOG

<http://summitblog.redhat.com/>

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

